Integrating Ethics in Computer Science with Inquiry-Based Learning

OER Reusable Abstract Model with Instance about Introductory Programming

This **Open Educational Resource (OER)** integrates the **ACM Code of Ethics and Professional Conduct** into the computer science curriculum using an **Inquiry-Based Learning (IBL)** approach. It provides both a generic, **reusable IBL model** and a specific **instantiation about programming control structures**.

The resource has three primary goals:

1. **Embed ethics directly** into software development practices.
2. **Develop critical thinking skills** by encouraging students to discuss, reflect, and articulate their own ethical positions. Specifically, students are expected to **generate questions**, **articulate their positions, self-evaluate**, and **avoid simple information regurgitation**.
3. Provide opportunities for students to **develop and practice essential soft skills**.

Explicitly designed for **Reuse, Revision, and Remixing**—core tenets of OER—the document is structured for maximum adaptability. It first presents a **generic, abstract IBL model**. This model incorporates **guided and open inquiry** and integrates the use of **generative AI (gen-AI) tools**. This generic model is then **instantiated with a sample case** to illustrate practical application and adaptation for creating additional lesson plans. Finally, the OER includes suggestions for **student data collection** with related instruments, as well as **commented partial solutions** illustrating the difference between average and excellent students' work.

Part of this material has already been formally **evaluated** (Rajan and Cortinovis, 2024). A larger evaluation trial is planned for **2026**, involving institutions in **Singapore, the US, and Italy**. An AI-based application is being developed to automatically generate additional instances.

1. Generic abstract model

The **lesson plan** is **grounded** in **Inquiry-Based Learning (IBL)**, encouraging students to **generate questions**, **construct explanations**, and **engage in self-reflection**. Following the **5E instructional model** (i.e., **Engage**, **Explore**, **Explain**, **Elaborate**, and **Evaluate**), this approach ensures a progressive and comprehensive learning path.

The **ACM Code of Ethics and Professional Conduct** is treated not as abstract theory but as a **living framework for professional decisions**.

The activity intentionally contrasts phases where **gen-AI use is restricted** (to foster independent **critical** reasoning) and phases where **gen-AI is permitted** (to highlight both its **potential** and its **limitations** in complex ethical reasoning). This dual stance helps students critically evaluate AI outputs rather than relying on them without scrutiny.

The activities for the first phase (Engage) and the final phase (Evaluate) are largely specified. The intermediate phases (Explore, Explain, and Elaborate) are outlined only in general terms and **will be detailed** according to the **specific computing topic** in which ethics will be integrated.

Phase	General purpose	Activity
Engage	Spark interest and curiosity. Situate learning in a meaningful context. Encourage questions and inquiry.	• **Preliminary familiarization.** Ask students to critically read the *ACM Code of Ethics and Professional Conduct*. Lead an introductory group discussion on the importance of ethics in computer science, prompting students to identify key issues and illustrate them with practical examples. • **Debate on ethical dilemmas.** The aim of this step is to introduce motivating and thought-provoking cases to illustrate that applying the *ACM Code of Ethics and Professional Conduct* is not always straightforward. A compelling example is decision-making in autonomous vehicles. Present the Moral Machine (Awad et al., 2018) scenarios to the class and, for each case, ask students to discuss the options, identify the most appropriate decision, justify their choice, and connect it to relevant aspects of the Code.
Explore	Provide direct experience of a phenomenon or concept. Encourage students to investigate their questions, test ideas, and solve problems.	• **First artifact development.** Ask students to create a first artifact (e.g., a short report, a conceptual diagram, or a prototype) on a specified topic, explicitly highlighting the connections with the *ACM Code of Ethics and Professional Conduct*. • **Documentation.** Students should submit the artifact.
Explain [guided inquiry]	Introduce conceptual tools to interpret evidence and construct explanations. Support multi-modal explanations and justify claims based on gathered evidence. Compare explanations generated by different students or groups.	• **Multi-modal extension of the artifact.** Ask students to improve the previous artifact or represent it using a different modality and set of tools, in order to address specified ethical implications. They may freely use Internet resources, but gen-AI tools are not allowed. • **Documentation.** Students should submit this new artifact.
Elaborate [open inquiry]	Apply concepts and explanations in new contexts to test general applicability. Reconstruct and extend understanding using multiple modes, such as written language, diagrams, graphics, and mathematics.	• **Further artifact enhancement.** Ask students to identify extensions or improvements to their previous artifact, integrating new ethical aspects that they must autonomously identify and justify. Here, they may freely use a gen-AI tool. To maximize ethical insight, students should formulate complex, analytical prompts that require the AI to engage with nuanced ethical reasoning, rather than generic queries. • **Documentation.** Students should submit the prompts they used with the tool, as well as the improved artifact.
Evaluate	Offer opportunities for students to review and reflect on their learning, understanding, and skills. Provide evidence of changes in students' knowledge, beliefs, and abilities.	• **Presentation of results.** Ask students to present their results to the class (5–10 minutes) followed by a 10-minute Q&A session. • **Peer interaction.** Require each student to ask at least one question to another student. • **Peer evaluation.** Ask students to evaluate one or more presenters on content, structure, and communication, noting first what they liked about the presentation and then what could be improved. • **Documentation.** Students should submit their presentation, questions, and evaluations.

 CC BY-NC-SA 4.0

2. Sample Instance: Ethics and Conditional Logic (Personalized Ads)

Adapted from *Personalized Ads Programming Assignment* by Natalie Garrett & Casey Fiesler (Internet Rules Lab)

This activity integrates the ACM Code of Ethics into the teaching of introductory control structures (if/elif/else statements) and basic data types in simple functions, using the context of designing a personalized advertising system. The focus is on how technical decisions in simple code directly impact ethical principles like Fairness, Privacy, and Harm.

It is an instance of the previous OER abstract model in the context of programming control structures.

Learning Goals

Upon completion, the student will be able to:

1. **Connect** specific conditional code structures to principles in the ACM Code (e.g., Fairness, Privacy, Harm Avoidance).
2. **Analyze** ethical trade-offs (e.g., maximizing revenue vs. protecting privacy) and make justified design choices.
3. **Appreciate** the utility and limitations of Generative AI (gen-AI) tools in complex ethical reasoning.
4. **Practice** essential soft skills in class debate (Engage), group discussion, and peer evaluation/presentation (Evaluate).

2.1 Lesson plan for the teacher

The activity centers on writing code to target ads based on user data, which carries inherent ethical risks (Fiesler, 2021). Here is an outline of the lesson plan, and the expected student focus:

Phase	Activity	Expected Student Inquiry/Focus
Engage (3 hours)	**Preliminary familiarization.** Ask students to critically read the ACM Code of Ethics and Professional Conduct. Lead an introductory group discussion on the importance of ethics in computer science, prompting students to identify key issues and illustrate them with practical examples. **Debate on ethical dilemmas.** Present the Moral Machine (Awad et al., 2018) scenarios to the class and, for each case, ask students to discuss the options, identify the most appropriate decision, justify their choice, and connect it to relevant aspects of the Code.	Introduce the idea that binary/conditional decisions (if/else logic) can have morally significant outcomes and highlight possible ethical conflicts.
Explore (2 hours)	**First artifact development.** Ask students to individually reflect on the ethical risks inherent in common ad personalization scenarios. Students should produce a brief report or a conceptual map on the following scenarios, identifying which ACM Code of Ethics principles are most at risk in each: 1. Pet Ownership: Targeting "pet insurance" ads based on a user having a large number of dog-related photos. 2. Social Isolation: Targeting ads for "self-improvement courses" based on a user having a very low friend count. 3. Wealth Inference: Targeting high-interest loan ads based on user location (zipcode) and age (to infer income). **Documentation**. Students should submit the document they produced.	Assess if students can articulate the link between the inferred data (e.g., low friend count) and the ethical risk (e.g., manipulation/harm).

Explain (2 hours)	**Multi-modal extension of the artifact.** Ask students to write Python functions (or use similar code with if/elif/else statements) to implement the ad targeting for Scenarios 1, 2, and 3 above. Students must provide comments on their code, explicitly linking their decision logic (the boolean expressions and the resulting ad content) to items in the ACM Code of Ethics. For example, justify the income thresholds in Scenario 3 with respect to Harm/Equity (Code 1.2). Gen-AI tools are NOT allowed. **Documentation.** Students should submit the code together with comments explaining the ethical considerations included.	Look for students using filters (e.g., blocking minors, avoiding predatory ads) as technical *ethical interventions*.
Elaborate (2 hours)	**Further artifact enhancement.** Ask students to create a Custom Ad Targeting function (Scenario 4) that could use sensitive data (e.g., political affiliation, relationship status, health interests). Students must use Gemini to identify the ethical trade-off between Privacy (Code 1.6) (avoiding sensitive data use) and Professional Competence (Code 2.1) (maximizing advertiser revenue). Students must formulate a complex, analytical prompt and refine their custom ad function or its accompanying documentation accordingly. **Documentation**. Students should submit the improved code, include comments linking the new extensions to the ACM Code, and document the prompts used with the gen-AI tool. The submitted documentation should also include a brief (2-3 sentence) justification for why the chosen prompt was designed to elicit complex ethical insight.	Look for students to identify the conflict and ask the AI for a technical compromise (e.g., data abstraction) rather than a simple ethical list. Assess the quality of the ethical reasoning and the effectiveness of the technical solution in balancing the trade-off.
Evaluate (3 hours)	**Presentation of results.** Ask students to present their results to the class (5–10 minutes) focused on their Custom Ad function logic (Scenario 4) and how they used the gen-AI tool to navigate the ethical trade-off. Followed by a 10-minute Q&A session. **Peer interaction.** Require each student to ask at least one question to another student. **Peer evaluation.** Ask students to evaluate one or more presenters on content, structure, and communication, noting first what they liked about the presentation and then what could be improved. **Documentation.** Students should submit their presentation, question(s), and evaluation(s).	

The activity is based on the following scenarios:

Scenario	**Data Used (Inferred)**	**Full Description of Inference and Risk (For Instructor Guidance)**
Scenario 1: Pet Insurance Ad	Pet ownership (dog/no dog)	**Inference:** Image recognition (e.g., from user photos) is a proxy for pet ownership. **Risk:** If the ML model used to classify images is biased (e.g., performs poorly on images of specific breeds or pets of minority owners), it risks **unfairly excluding** (Code 1.4) those users from seeing a relevant product.
Scenario 2: Self-Improvement Ad	Friend count (Social isolation)	**Inference:** A very low friend/follower count is used to **infer loneliness or social isolation/vulnerability**. **Risk:** Targeting individuals based on perceived vulnerability with manipulative advertising (e.g., promising immediate life changes) risks **psychological harm** (Code 1.2).

 CC BY-NC-SA 4.0

Scenario	Data Used (Inferred)	Full Description of Inference and Risk (For Instructor Guidance)
Scenario 3: Wealth Inference Ad	Zipcode and Age (financial status)	**Inference: Zipcodes** are strong predictors of **average neighborhood income and wealth**. Combining this with **Age** allows the system to infer a user's *likely* financial stability. **Risk:** This can lead to **digital redlining** (excluding ads for opportunities) or **predatory targeting** (pushing high-interest loans onto low-income groups), violating **Harm** (Code 1.2) and **Fairness** (Code 1.4).
Scenario 4: Custom Ad (Elaborate Phase)	Highly Sensitive PII (e.g., political affiliation, relationship status)	**Trade-off:** The system can maximize ad **Competence/Revenue** (Code 2.1) by using highly sensitive PII, but doing so drastically increases the risk to **Privacy** (Code 1.6) and **Harm** (manipulation) (Code 1.2). The solution requires a technical compromise.

 CC BY-NC-SA 4.0

2.2 Student's Instruction Sheet

This activity integrates the ACM Code of Ethics into the practice of writing simple functions with basic conditional logic (if/elif/else). You will analyze and implement ad targeting systems, focusing on the ethical implications of your code. You will be creating functions that will use user data to output personalized ads using conditionals.

The activities you are required to carry out, are described in the table below.

Phase	**Task**	**Deliverables**
Engage (3 hrs)	Preliminary Familiarization: Critically read the ACM Code of Ethics. Debate on Ethical Dilemmas: Participate in the discussion of Moral Machine scenarios.	Be prepared to discuss key principles in class. Identify which ACM principles are in conflict in a given scenario.
Explore (2 hrs)	Produce a Brief Report or Conceptual Map that addresses the following for each of scenarios 1, 2, and 3 (described in the following table): 1. Identify the single most relevant ACM Code principle at risk. 2. Provide a short justification (2-3 sentences) for *why* this principle is at risk.	Brief Report/Conceptual Map
Explain (2 hrs)	Write Python functions or similar code (if/elif/else) to implement the conditional logic for ad targeting in previous Scenarios 1, 2, and 3. Internet resources can be used, but gen-AI tools are not allowed.	Code file with explicit code comments linking your conditional logic to the ACM Code principles.
Elaborate (2 hrs)	Design a Custom Ad function for Scenario 4 (described in the following table) using the support of Gemini.	Refined Code, gen-AI Prompt, and 2-3 sentence justification for your analytical prompt design.
Evaluate (3 hrs)	Prepare and deliver a presentation on your Scenario 4 solution, followed by peer Q&A and evaluation.	Presentation, Written Question(s), and Peer Evaluation(s).

 CC BY-NC-SA 4.0

The table below describes the four scenarios and the program that decides on the ad to serve to a person on a social media platform. The program prompts the user for information and then returns text that describes ads based on their inputs.

Description	Personalized ad specifications
Scenario 1: For advertisers using the lowest tier of personalization on a social media platform, there is limited data about users available—we only know whether a user owns a dog, based on the analysis of the user's pictures.	Given a boolean returned by the function owns_dog, the ad program (function): • If the value of this boolean is True, prints text advertising dog food. The text must start with "Dog Food Ad: ". • Otherwise, it prints text advertising anything else. The text must start with "Ad: ".
Scenario 2*:* By combining data from a marketing platform with the results from personality tests that users have completed, we know that people with at least 500 Facebook friends are likely to be extroverts, and people with less than 500 friends are likely to be introverts. Based on analysis of user-uploaded photos using computer-vision algorithms, we also know that people who are extroverts are more likely to have many photos of dogs, and introverts are more likely to have many photos of cats. Market research has shown that dog people are highly likely to be positively influenced by advertisements for any product that includes dogs, and similar for cat people and cats.	Given how many friends a user has, the ad program (function): • If the user has at least (more than or equal to) 500 friends, prints ad text that includes a dog. The text must start with "Ad with Dog: ". • If the user has less than 500 friends, it prints ad text that includes a cat. The text must start with "Ad with Cat:".
Scenario 3: Based on ad clicks and third-party data shared through purchases on a "marketplace" platform, we have data about likely income averages for every zipcode in the country. We also know based on self-reported age how incomes vary based on life stage (e.g., college students have less disposable income no matter where they live). Advertisers strategically advertise to people who are more likely to be able to afford their product.	Given a user's *zipcode* and *age* as integers, the ad program (function): • If the *zipcode* is in Boulder (i.e., between and including 80301 and 80310) and the user is at least 25, prints ad text for an expensive product. The text must start with "Expensive Product Ad:". • If the *zipcode* is in Boulder and the user is younger than 25, prints ad text for an inexpensive product. The text must start with "Cheap Product Ad: ". • Otherwise, it prints ad text encouraging the user to visit Boulder. The text must start with "Tourist Ad: ".

 CC BY-NC-SA 4.0

Scenario 4*:* Now imagine that you have access to even more of a user's personal information based on their profile. In order to help advertisers get the most money out of their advertising packages, you tailor the ads a user will see based on personal details each user provides.	The ad program (function): Uses personal details from a social media profile to display at least 3 different custom ads depending on what the user has input. The only requirement is that the advertisements must start with "Custom Ad [number]:". Descriptions of the variables are below. The function must take at least 4 of the following values (as distinct parameters or as a single dictionary). • *name* – a string containing a user's name. (e.g. Avatar Aang) • *gender* – a string containing a user's gender (e.g. man) • *age* – an integer containing a user's age (e.g. 51) • *hometown* – a string containing a user's hometown (e.g. Southern Air Temple) • *relationship* – a string containing a user's relationship status (e.g. Crushing on Someone) • *birthday* – a string containing a user's birthday (e.g. August 21 12BG) • *number_of_friends* – an integer containing the number of friends a user has on social media (e.g. 243) • *political_affiliation* – a string containing a user's political affiliation (e.g. Independent)

 CC BY-NC-SA 4.0

2.3 Data Collection

This section suggests possible data to be collected, and the associated instruments.
These suggestions can be adapted based on specific research or instructional objectives.

Pre-Activity Survey (5 minutes)

1. **Familiarity**
    - How familiar are you with the ACM Code of Ethics? (1 = Not at all, 5 = Very familiar)
2. **Confidence**
    - How confident are you in applying ethical principles to technical design (e.g., database, software)? (1 = Not confident, 5 = Very confident)
3. **Expectations**
    - What do you expect to learn from this activity? *(open text)*

Post-Activity Survey (5–7 minutes)

1. **Learning Outcomes**
    - I feel more confident applying the ACM Code of Ethics to technical work. (1–5 Likert scale)
    - I can connect ethical principles to design choices in my field. (1–5 Likert scale)
2. **Engagement with IBL 5E**
    - Please rate how much each phase contributed to your learning (1=Least, 5=Most):
      Engage:
      Explore:
      Explain:
      Elaborate:
      Evaluate:
    - Which phase contributed most to your learning? Why? (open text)
3. **Use of gen-AI (Elaborate Phase)**
    - Using a gen-AI tool helped me broaden my ethical reflection. (1–5 Likert scale)
    - Please describe one useful or problematic aspect of using AI in this context. *(open text)*
4. **Global Perspectives**
    - Which principle of the ACM Code felt most relevant in your context? (open text, e.g., privacy, fairness, professional responsibility)

 CC BY-NC-SA 4.0

Artifact Evaluation Rubric (for Explore, Explain, Elaborate phases)

Criterion	Excellent (5)	Good (4)	Fair (3)	Limited (2)	Poor (1)
Clarity (technical accuracy & organization)	Artifact is complete, well-structured, and technically accurate	Mostly accurate and clear	Some inaccuracies or unclear structure	Many errors, weak organization	Incomplete / incoherent
Ethical Alignment (explicit reference to ACM Code)	Strong, explicit references to multiple, distinct ACM Code principles	Some explicit references	General mention without clear links	Very limited or vague references	No ethical consideration
Depth of Reasoning (quality of justification)	Thorough, insightful ethical reasoning with justification	Mostly well-justified, minor gaps	Some justification but superficial	Minimal reasoning, poorly justified	No reasoning
Innovation *(Elaborate only)* (integration of AI and new ethical insights)	AI prompts & outputs clearly integrated, expand ethical scope	AI used meaningfully, and adds some new insights	AI use adds little depth	AI use superficial or irrelevant	No AI use

---

Peer Evaluation Rubric (Evaluate phase)

Students rate a peer presentation (1–5 scale + comments):

1. **Content:** Was the presentation informative, accurate, and well-connected to the ACM Code?
2. **Structure:** Was it logically organized and easy to follow?
3. **Communication:** Was the delivery clear, engaging, and professional?
4. One thing I liked
5. One thing to improve

Reflection Prompts (Evaluate phase)

To be submitted individually:

1. Which ACM Code principle(s) most influenced your artifact design?
2. Did your understanding of the ACM Code change through this activity? How?
3. What did you learn from your peers' presentations about ethics in computing?

 CC BY-NC-SA 4.0

2.4 Commented Sample Solutions (with Code Snippets)

This section provides partial solutions illustrating the difference between average and excellent work.

2.4.1 Partial average solution provided by a student

This student focuses primarily on Privacy (Code 1.6). Their work is characterized by correct but minimal technical changes and a generic gen-AI prompt.

Explain Phase Artifact: Scenario 3 Code and Ethical Justification

The student correctly implements the conditional logic for ad targeting based on inferred wealth (zipcode and age) but with a simple structure.

Python Code Snippet (Scenario 3: Wealth Inference):

```
def target_wealth_ad(zipcode, age):
    # Ethical Link: Code 1.2 (Avoid Harm), Code 1.4 (Fairness)
    # Target high-value ads only in presumed wealthy zipcodes for users
    # over 25, who are statistically more stable.
    if 80301 <= zipcode <= 80310:  # Boulder, CO high-value area
        if age >= 25:
            return "Ad: High-End Investment Seminar"
        else:
            return "Ad: Basic Savings Account"
    elif age < 18:
        # Avoid targeting minors for financial products completely (Code 1.2).
        return "Ad: College Textbook Discounts"
    else:
        return "Ad: General Consumer Product" # Safe default
```

Student Justification (Low Complexity): "My main ethical concern is Privacy (Code 1.6) and Harm (Code 1.2). The code uses inferred income data, which is sensitive. I made sure to stick only to the required inputs (zipcode, age) and used a broad age filter (age < 18) to avoid targeting minors for financial harm. This is a basic way to follow Code 1.2."

Elaborate Phase Artifacts: Refined Custom Ad Function

The student uses a generic prompt, and their refinement is a simple justification for excluding sensitive data.

Python Code Snippet (Scenario 4: Custom Ad - Average):

```
def target_custom_ad_avg(user_data):
    # Available data in user_data: name, age, political_affiliation, relationship_status
    # Ethical decision: Only use name and age, exclude sensitive PII.
    name = user_data.get('name')
    age = user_data.get('age')

    # Justification (Code 1.6: Privacy, Code 1.4: Fairness):
    # Political affiliation is excluded from this logic to minimize privacy risk
    # and avoid unfair/manipulative targeting.
    if age >= 65:
        return f"Ad: Retirement Planning for {name}"
    else:
        return f"Ad: Travel Deals for {name}"
```

 CC BY-NC-SA 4.0

Student Justification for Refinement: "The AI suggested that Fairness (Code 1.4) is a big risk. My initial design for Scenario 4 *could* have used political affiliation, but the AI confirmed this violates Privacy (1.6) and risks Harm (1.2). My fix was to add a detailed comment block to my function explaining that sensitive PII like political_affiliation is deliberately excluded from the conditional logic".

2.4.2 Partial excellent solution provided by a student

This student immediately identifies a complex ethical trade-off between Professional Competence (Code 2.1) and Harm Avoidance/Privacy (Codes 1.2/1.6). Her solution uses data abstraction and time-control as technical ethical mechanisms.

Explain Phase Artifact: Scenario 3 Code and Ethical Justification

The student designs the conditional logic to filter out harmful possibilities even if it means sacrificing some revenue, demonstrating a higher-level structure.

Python Code Snippet (Scenario 3: Wealth Inference - Excellent):

```
def target_wealth_ad_excellent(zipcode, age, existing_debt=0):
    # Ethical Choice: Prioritize Harm Avoidance (Code 1.2) over Revenue (Code 2.1)

    # 1. Harm Filter: Block high-interest/predatory ads if user has high existing debt
    #    (even if inferred income is high).
    if existing_debt > 10000:
        return "Ad: Credit Counseling Service" # Ethical substitution

    # 2. Logic (Professional Competence - Code 2.1): Target based on inferred wealth
    if 80301 <= zipcode <= 80310 and age >= 35:
        # Highest confidence, target expensive product
        return "Ad: Luxury Home Equity Line"
    elif age >= 18:
        # Lower confidence/income inferred, target a safe financial product
        return "Ad: Basic Checking Account"
    else:
        # Safe default for minors/low-confidence users
        return "Ad: Educational Resources"
```

Student Justification (High Complexity): "My design focuses on the trade-off between Code 1.2 (Avoid Harm) and Code 2.1 (Competence). The code includes a Harm Filter that prioritizes avoiding predatory targeting (e.g., high-interest loans) if other data suggests high existing debt, even though an aggressive system (Code 2.1) would target them for maximum profit. This demonstrates a technical choice in the code's control structure to prioritize harm avoidance."

Elaborate Phase Artifacts: Refined Custom Ad Function

The student uses an analytical prompt and implements a technical solution based on data abstraction and time control.

Python Code Snippet (Scenario 4: Custom Ad - Excellent):

```
import datetime
def target_custom_ad_excellent(user_data):
    # Ethical Trade-off: Privacy (1.6) vs. Competence (2.1)
```

```
# Technical Solution 1: Time Control (Privacy/Competence)
# Check if the relationship_status data is older than 6 months.
last_update = user_data.get('status_last_updated', datetime.date.min)
if (datetime.date.today() - last_update).days > 180:
    return "Ad: General Interest Product" # Fallback if data is stale

# Technical Solution 2: Data Abstraction (Privacy/Fairness)
# Use a generic flag instead of the sensitive PII string.
is_seeking_social_product = user_data.get('relationship_status') in ['Single', 'Divorced']
if is_seeking_social_product:
    # Target a generalized social product (e.g., "Networking Events")
    return "Ad: Local Social Group App"
else:
    return "Ad: Home Improvement Tools"
```

Student Justification for Refinement: "The AI helped me confirm that using raw relationship status is a high ethical risk. My refinement uses two technical solutions: Time Control ensures we only use data when it's fresh enough to be relevant (satisfying Code 2.1) and Data Abstraction (is_seeking_social_product flag) allows me to target a product category without using the highly sensitive PII, mitigating the Privacy (1.6) risk while maintaining Competence (2.1)."

3. Acknowledgments

I would like to extend my sincere gratitude to Prof. Rajan Ranjidha for her contribution.

4. References

Awad, E., Dsouza, S., Kim, R., et al. (2018). The Moral Machine Experiment. *Nature, 563*(7729), 59–64. https://doi.org/10.1038/s41586-018-0637-6

Fiesler, C., Friske, M., Garrett, N., Muzny, F., Smith, J. J., & Zietz, J. (2021, March). Integrating Ethics Into Introductory Programming Classes. In *Proceedings of the 52nd ACM technical symposium on computer science education* (pp. 1027-1033).

Rajan, R., & Cortinovis, R. (2024). Ethical Integration In Computer Science Education: Leveraging Open Educational Resources and Generative Artificial Intelligence for Enhanced Learning. In *Proceedings of the 35th Annual Workshop of the Psychology of Programming Interest Group (PPIG), in conjunction with VL/HCC* (pp. 151–159). PPIG. https://www.ppig.org/files/2024-PPIG-35th-rajan.pdf

 CC BY-NC-SA 4.0